

Obfuscation Overhead for Single-Output Circuits

Proved for multi-output circuits; blocked on the NP-hardness of single-output circuit minimization

Status. Statement: AI-written from Zhenjian Lu, Noam Mazor, Igor C. Oliveira and Rafael Pass, *Lower Bounds on the Overhead of Indistinguishability Obfuscation*, IACR ePrint 2024/1524. Not yet checked by a human. Their main theorem gives an $\Omega(s/\log s)$ additive overhead lower bound for iO on *multi-output* circuits under $\text{NP} \not\subseteq \text{BPP}$. The single-output case is the first of their named “main open problems”, and the obstruction is precise: their route runs through the NP-hardness of Multi-MCSP, and the single-output analogue is the long-standing MCSP question.

Category. *Obfuscation.*

Conjecture (informal). *An obfuscator makes a program unintelligible; the question here is what that costs in size. If NP is easy there is no cost at all — an optimal obfuscator with zero overhead exists — so any lower bound needs a hardness assumption, and the weakest sensible one is that NP is hard. Under exactly that, the source shows an obfuscator for circuits with many output bits must add $\Omega(s/\log s)$ gates. Their argument goes through the hardness of deciding whether a multi-output function has a small circuit. For circuits with a single output bit the corresponding hardness question is one of the famous open problems of complexity theory, and the overhead question is open with it.*

1 The setting

For a circuit C of size s — the number of AND/OR/NOT gates — an indistinguishability obfuscator outputs a circuit of some size $\sigma(\lambda, s)$, and the *overhead* is how much larger that is than s . The question is how small the overhead can be.

The hardness assumption cannot be dropped: an optimal iO scheme with no circuit size overhead exists if $\text{NP} \subseteq \text{BPP}$. So $\text{NP} \not\subseteq$

BPP is the minimal assumption under which any lower bound can hold, and it is the one the source uses.

Theorem 1 (The multi-output bound, [LMOP24, Theorem 1]). *Under $\text{NP} \not\subseteq \text{BPP}$, there is no efficient indistinguishability obfuscation scheme for multi-output circuits $C : \{0,1\}^n \rightarrow \{0,1\}^n$ of size s that outputs circuits of size $s + o(s/\log s)$. Equivalently, to be secure, an efficient iO scheme must incur an $\Omega(s/\log s)$ additive overhead.*

The proof reinterprets a connection between obfuscation and meta-complexity due to Mazon and Pass — who showed that iO plus sub-exponentially secure one-way functions imply a gap version of MCSP is not NP-hard under Levin reductions — and turns it into a tool for proving impossibility. The missing ingredient is a meta-computational problem whose NP-hardness is known, and the source uses Multi-MCSP, NP-hard under randomized reductions by Ilango, Loff and Oliveira, strengthening it to a gap version under randomized approximate Levin reductions.

That choice is exactly what confines the result to multi-output circuits.

2 Notation and parameters

- s is the circuit size, counted as the number of AND/OR/NOT gates; λ the security parameter; $\sigma(\lambda, s)$ the output size of the obfuscator.
- A circuit is *multi-output* when it computes $C : \{0,1\}^n \rightarrow \{0,1\}^n$, and *single-output* when $C : \{0,1\}^n \rightarrow \{0,1\}$.
- MCSP is the minimum circuit size problem for single-output Boolean functions; *Multi-MCSP* its multi-output analogue.
- Overhead is *additive*: the target ruled out in Theorem 1 is $s + o(s/\log s)$.

3 The conjecture

Conjecture 1 ([LMOP24, §1.4]). *Theorem 1 holds for single-output Boolean circuits: under $\text{NP} \not\subseteq \text{BPP}$, there is no efficient indistinguishability obfuscation scheme for single-output circuits $C : \{0,1\}^n \rightarrow \{0,1\}$ of size s that outputs circuits of size $s + o(s/\log s)$.*

Remark 1 (How the source states it, and what is being supplied here). Page 12: “As our main open problems, we leave the extension of our (database-free) impossibility results to the setting of single-output circuits, the investigation of iO for multi-output circuits with output size $(1 + \Omega(1)) \cdot s$, and understanding the limits of our approach.” And earlier: “we currently do not know how to establish them for single-output Boolean circuits nor for multi-output circuits where we allow unbounded fan-in gates when measuring circuit size.”

The source names the direction without restating the bound, so the specific form of Conjecture 1 — the same $s + o(s/\log s)$ target, the same assumption — is this statement’s choice, and is the natural reading of “extension of our impossibility results”. A partial result at a weaker overhead is progress and should be reported as such.

Remark 2 (The obstruction, which is a named open problem). The route to Theorem 1 needs NP-hardness of a gap version of Multi-MCSP under randomized Levin reductions. For single-output circuits the corresponding statement is the NP-hardness of MCSP, which has been open for decades. The source is explicit that its own intermediate step is already hard: “In order to relax the iO assumption from Item 1 to circuits $C : \{0,1\}^n \rightarrow \{0,1\}^n$, it is sufficient to show in Item 2 above that Multi-MCSP retains its hardness on input functions $F : \{0,1\}^n \rightarrow \{0,1\}^n$. This remains a challenging open problem, and can be seen as an important step towards establishing the NP-hardness of MCSP for single-output Boolean functions.”

It also names the technical difficulty: the Multi-MCSP instances produced by the reduction from Set-Cover have the form $F : \{0,1\}^\ell \rightarrow \{0,1\}^m$ with m of order roughly 2^ℓ , and padding the input length to m would require printing a truth-table of size 2^m per coordinate — prohibitively large.

So a proof of Conjecture 1 by this route would carry a major complexity-theoretic result with it. That is the reason to expect it to be hard, and also the reason a route avoiding meta-complexity would be the interesting outcome.

Remark 3 (One padding obstacle the source did clear). Worth separating from the above, because it is easy to conflate. Going from arbitrary multi-output circuits $C : \{0,1\}^\ell \rightarrow \{0,1\}^m$ to the square

case $C : \{0,1\}^n \rightarrow \{0,1\}^n$ is solved: “in contrast with Multi-MCSP, the input of the iO procedure is a circuit instead of a truth-table. This allows us to establish through a simple padding argument that the hardness of iO for arbitrary multi-output circuits $C : \{0,1\}^\ell \rightarrow \{0,1\}^m$ implies its hardness for circuits of the form $C : \{0,1\}^n \rightarrow \{0,1\}^n$.” What remains open is the drop to a single output bit, not the shape of the multi-output case.

Remark 4 (The neighbouring parameter question, with its own named bottleneck). The second of the source’s main open problems is iO for multi-output circuits with output size $(1 + \Omega(1)) \cdot s$ — a constant multiplicative overhead rather than the additive $o(s/\log s)$ that Theorem 1 rules out. The source locates the bottleneck precisely: “The bottleneck lies in the proof of Lemma 20, i.e., on the gap between positive and negative instances of the NP-hardness of Multi-MCSP under Levin reductions. The choice of parameters governing this gap are constrained by the encoding argument presented near the end of the proof, which does not seem to allow an additive gap of the form $\Omega(s)$.” It also notes the overhead in its Theorems 1 and 2 can be slightly improved, from $\Omega(s/\log s)$ to $\Omega(s \log \log s / \log s)$. That is a different statement from Conjecture 1 and is not published here.

Remark 5 (What the source does prove for single-output circuits). Not nothing, but in a different model. It rules out iO for single-output *database-aided* circuits with an arbitrary polynomial overhead in circuit size, strengthening an impossibility of Goldwasser and Rothblum: where they considered circuits with access to an exponential-length database that the obfuscator accesses by oracle, the source’s result holds even for polynomial-size databases and even against obfuscators that may run in time polynomial in the database size and so read all of it. Conjecture 1 is the database-free case, which is why the source calls it “the extension of our (database-free) impossibility results”.

4 Bibliography

[LMOP24] Zhenjian Lu, Noam Mazor, Igor C. Oliveira and Rafael Pass. *Lower bounds on the overhead of indistinguishability obfuscation*. IACR ePrint 2024/1524. The source. The main theorem and the database-

aided result are in the abstract and Section 1; the remark that single-output circuits are not covered is on page 5; the Multi-MCSP padding difficulty and the single-output MCSP obstruction are on page 8; the list of main open problems is on page 12; the bottleneck remark is on page 33.

[MP24] Noam Mazon and Rafael Pass. *Gap MCSP is not (Levin) NP-complete in Obfuscopia*. Computational Complexity Conference (CCC) 2024. The connection between obfuscation and meta-complexity the source turns on its head to prove impossibility.

[ILO20] Rahul Ilango, Bruno Loff and Igor C. Oliveira. *NP-hardness of circuit minimization for multi-output functions*. Computational Complexity Conference (CCC) 2020. The NP-hardness of Multi-MCSP the source extends to a gap version under randomized approximate Levin reductions — and whose single-output analogue is the obstruction.

[GR07] Shafi Goldwasser and Guy N. Rothblum. *On best-possible obfuscation*. TCC 2007. The database-aided impossibility the source strengthens. [UNVERIFIED] — cited in the source as [GR07]; the bibliographic detail here is inferred.

[KMN+22] Cited by the source for the result that, assuming NP is hard, the existence of indistinguishability obfuscators for single-output circuits with polynomial output size implies one-way functions — which is what lets the source reduce its assumption to worst-case hardness. [UNVERIFIED] — cited in the source as [KMN+ 22]; neither the authors nor the venue were recovered during this harvest, and are deliberately not guessed.